



UNITED STATES PATENT AND TRADEMARK OFFICE

UNITED STATES DEPARTMENT OF COMMERCE
United States Patent and Trademark Office
Address: COMMISSIONER FOR PATENTS
P.O. Box 1450
Alexandria, Virginia 22313-1450
www.uspto.gov

APPLICATION NO.	FILING DATE	FIRST NAMED INVENTOR	ATTORNEY DOCKET NO.	CONFIRMATION NO.
09/650,800	08/30/2000	Paul S. Neuman	RA 5290(33012/289/101)	1186
27516	7590	09/17/2004	EXAMINER	
UNISYS CORPORATION MS 4773 PO BOX 64942 ST. PAUL, MN 55164-0942			VITAL, PIERRE M	
			ART UNIT	PAPER NUMBER
			2188	

DATE MAILED: 09/17/2004

Please find below and/or attached an Office communication concerning this application or proceeding.

Office Action Summary

Application No.

09/650,800

Applicant(s)

NEUMAN, PAUL S.

Examiner

Pierre M. Vital

Art Unit

2188

-- The MAILING DATE of this communication appears on the cover sheet with the correspondence address --
Period for Reply

A SHORTENED STATUTORY PERIOD FOR REPLY IS SET TO EXPIRE 3 MONTH(S) FROM THE MAILING DATE OF THIS COMMUNICATION.

- Extensions of time may be available under the provisions of 37 CFR 1.136(a). In no event, however, may a reply be timely filed after SIX (6) MONTHS from the mailing date of this communication.
- If the period for reply specified above is less than thirty (30) days, a reply within the statutory minimum of thirty (30) days will be considered timely.
- If NO period for reply is specified above, the maximum statutory period will apply and will expire SIX (6) MONTHS from the mailing date of this communication.
- Failure to reply within the set or extended period for reply will, by statute, cause the application to become ABANDONED (35 U.S.C. § 133). Any reply received by the Office later than three months after the mailing date of this communication, even if timely filed, may reduce any earned patent term adjustment. See 37 CFR 1.704(b).

Status

- 1) ☒ Responsive to communication(s) filed on 19 August 2004.
- 2a) ☒ This action is **FINAL**. 2b) ☐ This action is non-final.
- 3) ☐ Since this application is in condition for allowance except for formal matters, prosecution as to the merits is closed in accordance with the practice under *Ex parte Quayle*, 1935 C.D. 11, 453 O.G. 213.

Disposition of Claims

- 4) ☒ Claim(s) 1-25 is/are pending in the application.
- 4a) Of the above claim(s) _____ is/are withdrawn from consideration.
- 5) ☐ Claim(s) _____ is/are allowed.
- 6) ☒ Claim(s) 1-25 is/are rejected.
- 7) ☐ Claim(s) _____ is/are objected to.
- 8) ☐ Claim(s) _____ are subject to restriction and/or election requirement.

Application Papers

- 9) ☐ The specification is objected to by the Examiner.
- 10) ☒ The drawing(s) filed on 30 August 2000 is/are: a) ☒ accepted or b) ☐ objected to by the Examiner.
Applicant may not request that any objection to the drawing(s) be held in abeyance. See 37 CFR 1.85(a).
Replacement drawing sheet(s) including the correction is required if the drawing(s) is objected to. See 37 CFR 1.121(d).
- 11) ☐ The oath or declaration is objected to by the Examiner. Note the attached Office Action or form PTO-152.

Priority under 35 U.S.C. § 119

- 12) ☐ Acknowledgment is made of a claim for foreign priority under 35 U.S.C. § 119(a)-(d) or (f).
- a) ☐ All b) ☐ Some * c) ☐ None of:
- ☐ Certified copies of the priority documents have been received.
 - ☐ Certified copies of the priority documents have been received in Application No. _____.
 - ☐ Copies of the certified copies of the priority documents have been received in this National Stage application from the International Bureau (PCT Rule 17.2(a)).

* See the attached detailed Office action for a list of the certified copies not received.

Attachment(s)

- ☒ Notice of References Cited (PTO-892)
- ☐ Notice of Draftsperson's Patent Drawing Review (PTO-948)
- ☐ Information Disclosure Statement(s) (PTO-1449 or PTO/SB/08)
Paper No(s)/Mail Date _____
- ☐ Interview Summary (PTO-413)
Paper No(s)/Mail Date. _____
- ☐ Notice of Informal Patent Application (PTO-152)
- ☐ Other: _____

DETAILED ACTION

Response to Amendment

1. This Office Action is in response to applicant's communication filed August 19, 2004 in response to PTO Office Action mailed June 17, 2004. The Applicant's remarks and amendments to the claims and/or the specification were considered with the results that follow.
2. Claims 1-25 have been presented for examination in this application. In response to the last Office Action, no claims have been amended. No claims have been canceled or added. As a result, claims 1-25 are now pending in this application.

Response to Arguments

3. Applicant's arguments, see pages 6-7, filed August 19, 2004, with respect to the rejection(s) of claim(s) 1 under 35 USC 103(a) have been fully considered and are persuasive. Therefore, the rejection has been withdrawn. However, upon further consideration, a new ground(s) of rejection is made in view of Mayfield et al. (US6,446,167) and Lynch et al. (US6,061,766) and Lai (US5,564,035).
4. Applicant's arguments filed August 19, 2004 have been fully considered but they are not persuasive. As to the remarks, applicant asserted that:
 - (a) Lai does not teach or suggest the "first experiencing" and "second experiencing" steps.

Examiner respectfully traverses applicant's arguments for the following reasons. Examiner would like to emphasize that the L1 cache memory of Lai is local to the processor similar to the system claimed by applicant (see Figs. 2 and 3). Applicant's arguments that the Lai reference cannot have the "first and second experiencing steps" is clearly erroneous. Examiner would like to point out that Lai discloses that when a processor 200 is attempting to write to memory, then the processor 200 writes the data to the cache line of cache 202 that resulted in a cache hit (col. 1, lines 60-62). Because of the inclusion policy of the L2 cache 203, which dictates that all the contents of the L1 cache 202 are maintained in the L2 cache 203 (col. 2, lines 35-37); the L2 cache 203 will also experience a cache memory hit in response to a hit in L1 cache 202. Thus, it can be clearly seen that the system of Lai provides for a level one cache memory (i.e., L1 cache) hit and a level two cache memory (i.e., L2 cache) hit in response to a memory write as claimed by Applicant.

(b) There is no showing that Lai "records the location of data corresponding to read memory request" following a level one cache memory miss.

Examiner respectfully traverses applicant's arguments for the following reasons. Examiner would like to point out that the victim line recorded in Lai is a least recently used (LRU) line chosen when a primary cache (i.e., L1 cache) read miss occurs and stored in a victim buffer (col. 5, lines 22-30). Thus, by displacing the LRU line and storing the line in a buffer, it can be clearly seen that the address or location of the line is recorded as claimed by applicant.

- (c) The Examiner used mutually exclusive embodiments as taught by Lai.

Even though Lai discloses two different cache configurations, they are not mutually exclusive as asserted by Applicant. Note that Lai discusses using a partially inclusive cache as well (col. 6, lines 9-41). Because the prior art structure is capable of performing the intended use, then it meets the preferred embodiment of Lai.

- (d) Hazawa does not disclose a hierarchical cache memory structure.

Examiner would like to emphasize that Hazawa does teach a hierarchical cache structure as claimed by Applicant (see levels 0-3 of cache memory 3 in Fig. 2). It is to be noted that Applicant does not claim a vertical or horizontal cache memory. As such, applicant is arguing limitations, which are not recited in the rejected claims.

- (e) There is no showing of the generation of a parity error by a level two cache memory in Hazawa.

Note that the level 0 cache of Hazawa is equivalent to the level 1 cache of the present invention and the level 1 cache of Hazawa is equivalent to the level 2 cache of the present invention and so forth. Also note that Hazawa discloses an error has occurred at level 1 cache (col. 3, lines 45-57). As such, Hazawa teaches the generation of a parity error by a level two cache memory.

Claim Rejections - 35 USC § 102

5. The following is a quotation of the appropriate paragraphs of 35 U.S.C. 102 that form the basis for the rejections under this section made in this Office action:

A person shall be entitled to a patent unless –

(b) the invention was patented or described in a printed publication in this or a foreign country or in public use or on sale in this country, more than one year prior to the date of application for patent in the United States.

6. Claims 11, 14, 16 and 19 are rejected under 35 U.S.C. 102(b) as being anticipated by Lai (US5,564,035).

As per claim 11, Lai discloses a method of maintaining validity of data within a level one cache memory of a processor responsively coupled to a level two cache memory which is responsively coupled to a system memory bus [*L1 cache 202, L2 cache 203, system bus 208; Fig. 3*] comprising: formulating a write request [*processor 210 performs a write to memory; col. 2, lines 48-49*]; first experiencing a level one cache memory hit in response to said write memory request [*hit in L1 cache 202; col. 2, lines 48-52; col. 1, lines 60-62*]; second experiencing a level two cache memory hit in response to said first experiencing step [*L2 cache 203 detects a hit; col. 2, lines 48-49; col. 2, lines 35-37*]; and invalidating a portion of said level one cache memory corresponding to said write memory request in response to said second experiencing step [*hit in L1 cache corresponding cache line can be invalidated; col. 2, lines 52-57*].

As per claim 14, Lai discloses recording location of data in response to a level one cache read miss and a level two cache memory read miss to maintain multiprocessor coherency [*col. 6, line 45 – col. 7, line 11*].

As per claim 16, Lai discloses an apparatus comprising:

executing means for executing program instructions [*processor 200, 210; Fig. 3*];
level one caching means responsively coupled to said executing means for level one
caching data [*L1 cache 202; Fig. 3*]; accessing means responsively coupled to said
executing means and said level one caching means for accessing a data element if said
executing means requires accessing of said data element [*system bus 208; Fig. 3*]; level
two caching means responsively coupled to said requesting means for level two caching
data [*L2 cache 203; Fig. 3*]; and first invalidating means responsively coupled to said level
one caching means for invalidating said data element if said data element is a write data
element located within said level two caching means and within said level one caching
means [*hit in L2 cache and hit in L1 cache, corresponding cache line can be invalidated; col. 2, lines
52-57*].

As per claim 19, Lai discloses recording location of data in response to a level
one cache read miss and a level two cache memory read miss to maintain
multiprocessor coherency [*col. 6, line 45 – col. 7, line 11*].

Claim Rejections - 35 USC § 103

7. The following is a quotation of 35 U.S.C. 103(a) which forms the basis for all obviousness rejections set forth in this Office action:

(a) A patent may not be obtained though the invention is not identically disclosed or described as set forth in section 102 of this title, if the differences between the subject matter sought to be patented and the prior art are such that the subject matter as a whole would have been obvious at the time the invention was made to a person having ordinary skill in the art to which said subject matter pertains. Patentability shall not be negated by the manner in which the invention was made.

8. Claims 1-4 are rejected under 35 U.S.C. 103(a) as being unpatentable over Mayfield et al. (US6,446,167) and Lynch et al. (US6,061,766) and Lai (US5,564,035).

As per claim 1, Mayfield discloses a data processing system having a system bus and having a processor [*fabric 204; processors 101,103; Fig. 1*], with a level one cache memory [*L1 caches 112, 113; Fig. 1*], responsively coupled to a level two cache memory [*L2 cache 118; Fig. 1*], which is responsively coupled to a level three cache memory which is directly coupled to at least one memory storage [*L3 cache 205 directly coupled to main memory 102; Fig. 1*].

However, Mayfield does not specifically teach a circuit for Snooping said system bus and a first logic which invalidates a corresponding level one cache memory location in response to either a non-local write as recited in the claim.

Lynch discloses a pipelined snoop bus for maintaining coherence among caches in a multiprocessor configuration (col. 1, lines 9-11).

Since the technology for implementing a circuit for snooping a system bus was well known and since snooping a system bus provides a snoop process for ensuring cache coherency, an artisan would have been motivated to implement a circuit for snooping a system bus in the system of Mayfield. Thus, it would have been obvious to

one of ordinary skill in the art, at the time the invention was made, to modify the system of Mayfield to include a circuit for snooping a system bus because it was well known to provide a snoop process for ensuring cache coherency as taught by Lynch.

Lai discloses invalidating a corresponding level one cache memory location in response to a non-local write [*when second processor performs write to memory, hit in L1 cache, corresponding cache line can be invalidated*; col. 2, lines 41-57] to maintain multiprocessor coherency. Since the technology for implementing invalidating a corresponding level one cache memory location in response to a non-local write was well known and since invalidating a corresponding level one cache memory location in response to a non-local write maintains multiprocessor coherency, an artisan would have been motivated to implement invalidating a corresponding level one cache memory location in response to a non-local write in the system of Mayfield and Lynch. Thus, It would have been obvious to one of ordinary skill in the art, at the time the invention was made, to modify the system of Mayfield and Lynch to include invalidating a corresponding level one cache memory location in response to a non-local write because it was well known to maintain multiprocessor coherency as taught by Lai.

As per claim 2, the combination of Mayfield and Lai does not specifically teach a second logic which inhibits said first logic from invalidating for mode 3 requests without ownership as recited in the claim.

Lynch discloses a second logic which inhibits said first logic from invalidating for mode 3 requests without ownership [*snoop requests checks for the presence of an object in cache; only requests for exclusive use which match cache tags are invalidated; there is no invalidate done when there is no hit*; Fig. 4; col. 4, lines 19-30; *it is clearly obvious that any computer system uses a combination of logic to produce output based on the rules of logic it is designed to follow; clearly, the use of multiple logics is an inherent feature of any computer system*] to provide a snoop process for ensuring cache coherency (col. 2, lines 34).

Since the technology for implementing inhibiting a logic from invalidating mode 3 requests without ownership was well known and since inhibiting invalidating mode 3 requests without ownership provides a snoop process for ensuring cache coherency, an artisan would have been motivated to implement inhibiting invalidating mode 3 requests without ownership in the system of Mayfield and Lai. Thus, it would have been obvious to one of ordinary skill in the art, having the teachings of Mayfield and Lai and Lynch before him at the time the invention was made, to modify the system of Mayfield and Lai to include implementing inhibiting a logic from invalidating mode 3 requests without ownership because it was well known to provide a snoop process for ensuring cache coherency as taught by Lynch.

As per claim 3, the combination of Mayfield and Lai does not specifically teach third logic which invalidates said corresponding cache memory location in response to a SNOOP hit as recited in the claim.

Lynch discloses invalidating a corresponding cache memory location in response to a SNOOP hit [Fig. 4; col. 4, lines 24-26; *it is clearly obvious that any computer system uses a combination of logic to produce output based on the rules of logic it is designed to follow; clearly, the use of multiple logics is an inherent feature of any computer system*] to provide a snoop process for ensuring cache coherency (col. 2, lines 34).

Since the technology for implementing invalidating a corresponding cache memory location in response to a SNOOP hit was well known and since invalidating a corresponding cache memory location in response to a SNOOP hit provides a snoop process for ensuring cache coherency, an artisan would have been motivated to implement invalidating a corresponding cache memory location in response to a SNOOP hit in the system of Mayfield and Lai. Thus, it would have been obvious to one of ordinary skill in the art, having the teachings of Mayfield and Lai and Lynch before him at the time the invention was made, to modify the system of Mayfield and Lai to include invalidating a corresponding cache memory location in response to a SNOOP hit because it was well known to provide a snoop process for ensuring cache coherency as taught by Lynch.

As per claim 4, Mayfield does not specifically teach fourth logic which records location of data in response to a level one cache read miss and a level two cache memory read miss as recited in the claim.

Lai discloses recording location of data in response to a level one cache read miss and a level two cache memory read miss to maintain multiprocessor coherency

Art Unit: 2188

(col. 6, line 45 – col. 7, line 11). Since the technology for implementing recording location of data in response to a level one cache read miss and a level two cache memory read miss was well known and since recording location of data in response to a level one cache read miss and a level two cache memory read miss maintains multiprocessor coherency, an artisan would have been motivated to implement recording location of data in response to a level one cache read miss and a level two cache memory read miss in the system of Mayfield. Thus, It would have been obvious to one of ordinary skill in the art, at the time the invention was made, to modify the system of Mayfield to include recording location of data in response to a level one cache read miss and a level two cache memory read miss because it was well known to maintain multiprocessor coherency as taught by Lai.

9. Claim 5 is rejected under 35 U.S.C. 103(a) as being unpatentable over Mayfield et al. (US6,446,167) and Lynch et al. (US6,061,766) and Lai (US5,564,035) and Hazawa (US4,891,809).

As per claim 5, the combination of Mayfield and Lai and Lynch discloses the claimed invention as detailed above in the previous paragraphs. However, Mayfield and Lai and Lynch do not specifically teach invalidating a level one cache memory in response to a level two cache memory generating a parity error as recited in the claim.

Hazawa discloses invalidating a level one cache memory in response to a level two cache memory generating a parity error [col.3, lines 38-48].

It would have been obvious to one of ordinary skill in the art, having the teachings of Mayfield and Lynch and Hazawa before him at the time the invention was made, to modify the system of Mayfield and Lynch to include invalidating a level one cache memory in response to a level two cache memory generating a parity error because it was well known to provide a cache memory having a normal error checking mode as taught by Hazawa by the provision of a sequential verification logic circuit for generating error indicating signals in sequence within the cache memory unit [col. 1, lines 30-37] as taught y Hazawa.

10. Claim 6 is rejected under 35 U.S.C. 103(a) as being unpatentable over Mayfield et al. (US6,446,167) and Lai (US5,564,035).

As per claim 6, Mayfield discloses a data processing system comprising a level one cache memory [*L1 caches 112, 113*; Fig. 1]; a level two cache memory responsively coupled to said level one cache memory [*L2 cache 118*; Fig. 1]; a system bus [*fabric 204*; Fig. 1]; a memory storage unit [*main memory*; col. 4, lines 19-22]; a level three memory responsively coupled to said level two cache memory via said system bus and

Art Unit: 2188

responsively coupled to said memory storage unit [*L3 cache 205 directly coupled to main memory 102; Fig. 1*].

However, Mayfield does not specifically teach a first circuit to invalidate a corresponding portion of said level one cache memory in response to a level one cache memory write hit and a level two cache memory hit as recited in the claim.

Lai discloses a first circuit to invalidate a corresponding portion of said level one cache memory in response to a level one cache memory write hit and a level two cache memory hit [*hit in L2 cache and hit in L1 cache, corresponding cache line can be invalidated; col. 2, lines 41-57*] to maintain multiprocessor coherency.

Since the technology for implementing invalidating a corresponding level one cache memory location in response to a level one write hit and a level two write hit was well known and since invalidating a corresponding level one cache memory location in response to a level one write hit and a level two write hit maintains multiprocessor coherency, an artisan would have been motivated to implement invalidating a corresponding level one cache memory location in response to a level one write hit and a level two write hit in the system of Mayfield. Thus, It would have been obvious to one of ordinary skill in the art, at the time the invention was made, to modify the system of Mayfield to include invalidating a corresponding level one cache memory location in response to a level one write hit and a level two write hit because it was well known to maintain multiprocessor coherency as taught by Lai.

As per claim 9, Mayfield does not specifically teach fourth logic which records location of data in response to a level one cache read miss and a level two cache memory read miss as recited in the claim.

Lai discloses recording location of data in response to a level one cache read miss and a level two cache memory read miss to maintain multiprocessor coherency (col. 6, line 45 – col. 7, line 11). Since the technology for implementing recording location of data in response to a level one cache read miss and a level two cache memory read miss was well known and since recording location of data in response to a level one cache read miss and a level two cache memory read miss maintains multiprocessor coherency, an artisan would have been motivated to implement recording location of data in response to a level one cache read miss and a level two cache memory read miss in the system of Mayfield. Thus, It would have been obvious to one of ordinary skill in the art, at the time the invention was made, to modify the system of Mayfield to include recording location of data in response to a level one cache read miss and a level two cache memory read miss because it was well known to maintain multiprocessor coherency as taught by Lai.

11. Claims 7 and 8 are rejected under 35 U.S.C. 103(a) as being unpatentable over Mayfield et al. (US6,446,167) and Lai (US5,564,035) and Lynch et al. (US6,061,766).

As per claim 7, the combination of Mayfield and Lai does not specifically teach a second logic which inhibits said first logic from invalidating for mode 3 requests without ownership as recited in the claim.

Lynch discloses a second logic which inhibits said first logic from invalidating for mode 3 requests without ownership [*snoop requests checks for the presence of an object in cache; only requests for exclusive use which match cache tags are invalidated; there is no invalidate done when there is no hit; Fig. 4; col. 4, lines 19-30; it is clearly obvious that any computer system uses a combination of logic to produce output based on the rules of logic it is designed to follow; clearly, the use of multiple logics is an inherent feature of any computer system*] to provide a snoop process for ensuring cache coherency (col. 2, lines 34).

Since the technology for implementing inhibiting a logic from invalidating mode 3 requests without ownership was well known and since inhibiting invalidating mode 3 requests without ownership provides a snoop process for ensuring cache coherency, an artisan would have been motivated to implement inhibiting invalidating mode 3 requests without ownership in the system of Mayfield and Lai. Thus, it would have been obvious to one of ordinary skill in the art, having the teachings of Mayfield and Lai and Lynch before him at the time the invention was made, to modify the system of Mayfield and Lai to include implementing inhibiting a logic from invalidating mode 3 requests without

Art Unit: 2188

ownership because it was well known to provide a snoop process for ensuring cache coherency as taught by Lynch.

As per claim 8, the combination of Mayfield and Lai does not specifically teach third logic which invalidates said corresponding cache memory location in response to a SNOOP hit as recited in the claim.

Lynch discloses invalidating a corresponding cache memory location in response to a SNOOP hit [Fig. 4; col. 4, lines 24-26; *it is clearly obvious that any computer system uses a combination of logic to produce output based on the rules of logic it is designed to follow; clearly, the use of multiple logics is an inherent feature of any computer system*] to provide a snoop process for ensuring cache coherency (col. 2, lines 34).

Since the technology for implementing invalidating a corresponding cache memory location in response to a SNOOP hit was well known and since invalidating a corresponding cache memory location in response to a SNOOP hit provides a snoop process for ensuring cache coherency, an artisan would have been motivated to implement invalidating a corresponding cache memory location in response to a SNOOP hit in the system of Mayfield and Lai. Thus, it would have been obvious to one of ordinary skill in the art, having the teachings of Mayfield and Lai and Lynch before him at the time the invention was made, to modify the system of Mayfield and Lai to include invalidating a corresponding cache memory location in response to a SNOOP hit because it was well known to provide a snoop process for ensuring cache coherency as taught by Lynch.

12. Claim 10 is rejected under 35 U.S.C. 103(a) as being unpatentable over Mayfield et al. (US6,446,167) and Lai (US5,564,035) and Hazawa (US4,891,809).

As per claim 10, the combination of Mayfield and Lai and Lynch discloses the claimed invention as detailed above in the previous paragraphs. However, Mayfield and Lai and Lynch do not specifically teach invalidating a level one cache memory in response to a level two cache memory generating a parity error as recited in the claim.

Hazawa discloses invalidating a level one cache memory in response to a level two cache memory generating a parity error [col.3, lines 38-48].

It would have been obvious to one of ordinary skill in the art, having the teachings of Mayfield and Lynch and Hazawa before him at the time the invention was made, to modify the system of Mayfield and Lynch to include invalidating a level one cache memory in response to a level two cache memory generating a parity error because it was well known to provide a cache memory having a normal error checking mode as taught by Hazawa by the provision of a sequential verification logic circuit for generating error indicating signals in sequence within the cache memory unit [col. 1, lines 30-37] as taught y Hazawa.

13. Claims 12-13 and 17-18 are rejected under 35 U.S.C. 103(a) as being unpatentable over Lai (US5,564,035) and Lynch et al. (US6,061,766).

As per claims 12 and 17, Lai does not specifically teach a second logic which inhibits said first logic from invalidating for mode 3 requests without ownership as recited in the claim.

Lynch discloses a second logic which inhibits said first logic from invalidating for mode 3 requests without ownership [*snoop requests checks for the presence of an object in cache; only requests for exclusive use which match cache tags are invalidated; there is no invalidate done when there is no hit; Fig. 4; col. 4, lines 19-30; it is clearly obvious that any computer system uses a combination of logic to produce output based on the rules of logic it is designed to follow; clearly, the use of multiple logics is an inherent feature of any computer system*] to provide a snoop process for ensuring cache coherency (col. 2, lines 34).

Since the technology for implementing inhibiting a logic from invalidating mode 3 requests without ownership was well known and since inhibiting invalidating mode 3 requests without ownership provides a snoop process for ensuring cache coherency, an artisan would have been motivated to implement inhibiting invalidating mode 3 requests without ownership in the system of Lai. Thus, it would have been obvious to one of ordinary skill in the art, having the teachings of Lai and Lynch before him at the time the invention was made, to modify the system of Lai to include implementing inhibiting a

logic from invalidating mode 3 requests without ownership because it was well known to provide a snoop process for ensuring cache coherency as taught by Lynch.

As per claims 13 and 18, the combination of Lai does not specifically teach third logic which invalidates said corresponding cache memory location in response to a SNOOP hit as recited in the claim.

Lynch discloses invalidating a corresponding cache memory location in response to a SNOOP hit [Fig. 4; col. 4, lines 24-26; *it is clearly obvious that any computer system uses a combination of logic to produce output based on the rules of logic it is designed to follow; clearly, the use of multiple logics is an inherent feature of any computer system*] to provide a snoop process for ensuring cache coherency (col. 2, lines 34).

Since the technology for implementing invalidating a corresponding cache memory location in response to a SNOOP hit was well known and since invalidating a corresponding cache memory location in response to a SNOOP hit provides a snoop process for ensuring cache coherency, an artisan would have been motivated to implement invalidating a corresponding cache memory location in response to a SNOOP hit in the system of Lai. Thus, it would have been obvious to one of ordinary skill in the art, having the teachings of Lai and Lynch before him at the time the invention was made, to modify the system of Lai to include invalidating a corresponding cache memory location in response to a SNOOP hit because it was well known to provide a snoop process for ensuring cache coherency as taught by Lynch.

14. Claims 15 and 20 are rejected under 35 U.S.C. 103(a) as being unpatentable over Lai (US5,564,035) and Hazawa (US4,891,809).

As per claims 15 and 20, the combination of Lai and Lynch discloses the claimed invention as detailed above in the previous paragraphs. However, Lai and Lynch do not specifically teach invalidating a level one cache memory in response to a level two cache memory generating a parity error as recited in the claim.

Hazawa discloses invalidating a level one cache memory in response to a level two cache memory generating a parity error [col.3, lines 38-48].

It would have been obvious to one of ordinary skill in the art, having the teachings of Lai and Hazawa before him at the time the invention was made, to modify the system of Lai to include invalidating a level one cache memory in response to a level two cache memory generating a parity error because it was well known to provide a cache memory having a normal error checking mode as taught by Hazawa by the provision of a sequential verification logic circuit for generating error indicating signals in sequence within the cache memory unit [col. 1, lines 30-37] as taught y Hazawa.

15. Claim 21 is rejected under 35 U.S.C. 103(a) as being unpatentable over Mayfield et al. (US6,446,167) and Hazawa (US4,891,809).

As per claim 21, Mayfield discloses an instruction processor [*processors 101, 103; Fig. 1*]; a level one cache memory directly coupled to said instruction processor [*L1 caches 112, 113; Fig. 1*]; a level two cache memory directly coupled to said level one cache memory [*L2 cache 118; Fig. 1*].

However, Mayfield does not specifically teach a data element having a parity error stored in said level two cache memory and a facility responsively coupled to said level one cache memory and said level two cache memory which detects said parity error of said data element and invalidates a corresponding data element within said level two cache memory as recited in the claim.

Hazawa discloses invalidating data in a level two cache memory in response to a parity error of a data element to provide a cache memory with an error checking mode [*col.3, lines 38-51; col. 3, lines 38-59*].

It would have been obvious to one of ordinary skill in the art, having the teachings of Mayfield and Hazawa before him at the time the invention was made, to modify the system of Mayfield to include invalidating data in a level two cache memory

in response to a parity error of a data element because it was well known to provide a cache memory having a normal error checking mode as taught by Hazawa.

16. Claim 22 is rejected under 35 U.S.C. 103(a) as being unpatentable over Mayfield et al. (US6,446,167) and Hazawa (US4,891,809) and Lynch et al (US6,061,766).

As per claim 22, the combination of Mayfield and Hazawa discloses the claimed invention as detailed above in the previous paragraphs. However, Mayfield and Hazawa do not specifically teach a level one cache memory comprising a level one instruction cache memory and a level one operand cache memory as recited in the claim.

Lynch discloses a level one cache comprising a data cache for storing data as it is passed back and forth from the execution units of the processor and an instruction cache holding instructions prior to execution by the processor's execution units (col. 3, lines 43-48).

It would have been obvious to one of ordinary skill in the art, having the teachings of Mayfield and Hazawa and Lynch before him at the time the invention was made, to modify the system of Mayfield and Hazawa to include a level one instruction cache memory and a level one operand cache memory because a level one data cache was well known for storing data as it is passed back and forth from the execution units of the processor and a level one instruction cache was well known for holding instructions prior to execution by the processor's execution units as taught by Lynch.

17. Claims 23-25 are rejected under 35 U.S.C. 103(a) as being unpatentable over Mayfield et al. (US6,446,167) and Hazawa (US4,891,809) and Lynch et al (US6,061,766) and Lai (US5,564,035).

As per claim 23, the combination of Mayfield and Hazawa and Lynch discloses the claimed invention as detailed above in the previous paragraphs. However, Mayfield and Hazawa and Lynch do not specifically teach invalidating a write data element of a level one cache memory in response to a level one cache memory write hit and a level two cache memory hit as recited in the claim.

Lai discloses invalidating a write data element of a level one cache memory in response to a level one cache memory write hit and a level two cache memory hit [*hit in L2 cache and hit in L1 cache, corresponding cache line can be invalidated*; col. 2, lines 41-57] to maintain multiprocessor coherency.

Since the technology for implementing invalidating a corresponding level one cache memory location in response to a level one write hit and a level two write hit was well known and since invalidating a corresponding level one cache memory location in response to a level one write hit and a level two write hit maintains multiprocessor coherency, an artisan would have been motivated to implement invalidating a

Art Unit: 2188

corresponding level one cache memory location in response to a level one write hit and a level two write hit in the system of Mayfield and Hazawa. Thus, It would have been obvious to one of ordinary skill in the art, at the time the invention was made, to modify the system of Mayfield and Hazawa to include invalidating a corresponding level one cache memory location in response to a level one write hit and a level two write hit because it was well known to maintain multiprocessor coherency as taught by Lai.

As per claim 24, Lynch discloses a snooping circuit [*snoop queue 402*; Fig. 4].

As per claim 25, Lynch discloses said write data element is located within said level one operand cache memory [col. 3, lines 43-45].

Conclusion

18. The prior art made of record and not relied upon is considered pertinent to applicant's disclosure. Applicant is required under 37 C.F.R. § 1.111 (c) to consider these references fully when responding to this action. The documents cited therein teach a level two cache coupled to a level three cache which is directly coupled to a memory unit through a system bus.

19. Applicant's amendment necessitated the new ground(s) of rejection presented in this Office action. Accordingly, **THIS ACTION IS MADE FINAL**. See MPEP § 706.07(a). Applicant is reminded of the extension of time policy as set forth in 37 CFR 1.136(a).

A shortened statutory period for reply to this final action is set to expire THREE MONTHS from the mailing date of this action. In the event a first reply is filed within TWO MONTHS of the mailing date of this final action and the advisory action is not mailed until after the end of the THREE-MONTH shortened statutory period, then the shortened statutory period will expire on the date the advisory action is mailed, and any extension fee pursuant to 37 CFR 1.136(a) will be calculated from the mailing date of the advisory action. In no event, however, will the statutory period for reply expire later than SIX MONTHS from the date of this final action.

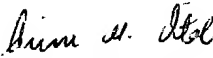
20. Any inquiry concerning this communication or earlier communications from the examiner should be directed to Pierre M. Vital whose telephone number is (703) 306-5839. The examiner can normally be reached on Mon-Fri, 8:30 am - 6:00 pm, alternate Friday off.

If attempts to reach the examiner by telephone are unsuccessful, the examiner's supervisor, Mano Padmanabhan can be reached on (703) 306-2903. The fax phone number for the organization where this application or proceeding is assigned is 703-872-9306.

Art Unit: 2188

Information regarding the status of an application may be obtained from the Patent Application Information Retrieval (PAIR) system. Status information for published applications may be obtained from either Private PAIR or Public PAIR. Status information for unpublished applications is available through Private PAIR only. For more information about the PAIR system, see <http://pair-direct.uspto.gov>. Should you have questions on access to the Private PAIR system, contact the Electronic Business Center (EBC) at 866-217-9197 (toll-free).

September 15, 2004


Pierre M. Vital
Examiner
Art Unit 2188